

SANDIA REPORT

SAND96-1916 • UC-903

Unlimited Release


Printed August 1996

The Waveform Correlation Event Detection System Project, Phase I: Issues in Prototype Development and Testing

Chris Young, Mark Harris, Judy Beiriger, Susan Moore, Julian Trujillo,
Mitch Withers, Richard Aster

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550
for the United States Department of Energy
under Contract DE-AC04-94AL85000

Approved for public release; distribution is unlimited.



SF2900Q(8-81)

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

Available to the public from
National Technical Information Service
US Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01

SAND96-1916
Unlimited Release
Printed August 1996

Distribution
Category UC-903

The Waveform Correlation Event Detection System Project, Phase I: Issues in Prototype Development and Testing

Chris Young, Geophysics Department
Mark Harris, Cooperative Monitoring Department
Judy Beiriger, Susan Moore, Julian Trujillo, Decision Support Systems Department
Sandia National Laboratories, Albuquerque, NM 87185

Mitch Withers and Richard Aster
Department of Earth and Atmospheric Sciences
New Mexico Institute of Mining and Technology
Socorro, NM 87801

Abstract

A study by Shearer (1994) using long-period seismic data showed that seismic events can be detected and located based on correlations of processed waveform profiles with the profile expected for an event. In this technique both time and space are discretized and events are found by forming profiles and calculating correlations for all time-distance points. Events are declared at points with large correlations. In the first phase of the Waveform Correlation Event Detection System (WCEDS) Project at Sandia Labs we have developed a prototype automatic event detection system based on Shearer's work which shows promise for treaty monitoring applications. Many modifications have been made to meet the requirements of the monitoring environment. A new full matrix multiplication has been developed which can reduce the number of computations needed for the data correlation by as much as two orders of magnitude for large grids. New methodology has also been developed to deal with the problems caused by false correlations (sidelobes) generated during the correlation process. When an event has been detected, masking matrices are set up which will mask all correlation sidelobes due to the event, allowing other events with intermingled phases to be found. This process is repeated until a detection threshold is reached. The system was tested on one hour of Incorporated Research Institutions for Seismology (IRIS) broadband data and built all 4 of the events listed in the National Earthquake Information Center (NEIC) Preliminary Determination of Epicenters (PDE) which were observable by the IRIS network. A continuous execution scheme has been developed for the system but has not yet been implemented. Improvements to the efficiency of the code are in various stages of development. Many refinements would have to be made to the system before it could be used as part of an actual monitoring system, but at this stage we know of no clear barriers which would prevent an eventual implementation of the system.

Acknowledgments

Helpful reviews of this report were provided by Marianne Walck and Eric Chael. This work was funded by the Automatic Data Processing Program at Sandia National Laboratories under Department of Energy contract DOE ST485D.

Contents

Abstract.....	i
Acknowledgments.....	ii
List of Tables.....	iv
List of Figures.....	v
Introduction	1
Event detection and location	1
Automatic event detection systems.....	1
Waveform correlation event detection.....	2
Development and Testing of the WCEDS Prototype	3
Waveform pre-processing.....	3
Master images	6
Event detection.....	10
The correlation matrix	15
False correlations	17
Event stripping and event masking.....	20
Continuous execution.....	28
Discussion	30
Conclusions	32
References	34
Distribution.....	35

List of Tables

Table 1.	PDE events for October 2, 1993 08:30:00 to 09:30:00	17
-----------------	--	-----------

List of Figures

Figure 1.	The basic waveform correlation event detection scheme	4
Figure 2.	Data pre-processing and waveform correlation.....	6
Figure 3.	Empirical master images.....	8
Figure 4.	Travel time-based master images	9
Figure 5.	Master image correlation widths and distance discretization...	11
Figure 6.	Correlation redundancy and a new detection algorithm	14
Figure 7.	C matrix examples	16
Figure 8.	The generation of false correlations.....	19
Figure 9.	Event stripping	21
Figure 10.	Use of the X matrix	23
Figure 11.	Determining the contributing phases	24
Figure 12.	Pre-masking and post-masking detector output	26
Figure 13.	WCEDS vs. PDE	27
Figure 14.	Continuous execution.....	29

Introduction

To support compliance with a global Comprehensive Test Ban Treaty (CTBT) it is desirable to achieve the lowest possible seismic event detection thresholds. The automatic seismic event detection systems currently used by the verification community represent vast improvements over the previous systems and have undoubtedly helped to significantly lower detection thresholds, but there is still ample room for improvement. Recent GSETT3 (Group of Scientific Experts Technical Test #3) results from the IDC (International Data Center) show that analysts reject 50% of the automatically detected events and manually build 20% of the events in the final bulletin (Bob North, personal comm., March 1996). While these figures can be expected to improve somewhat as the IDC software is tuned and refined, it is not clear that the current generation of systems can ever achieve significantly better results. In this paper we will discuss the development and testing of a prototype of a new type of automatic seismic event detection system which could provide substantial improvements and thereby significantly improve CTBT monitoring capabilities.

Event detection and location

While the process of event location has been revolutionized by the advent of computers, the process of event detection has remained at least in part stubbornly resistant to automation. Even the most sophisticated, highly automated systems in use today require review by analysts of their output to ensure bulletin reliability. Understanding the process whereby an analyst deems an event valid may be the key to building a better automated detection system. Humans have a tremendous capacity to recognize complex patterns in the data, especially with regard to subtle changes in frequency content. To date no automated detection system has been able to completely match this performance. It is not uncommon for an experienced analyst to be able to detect and locate an event recorded by a network based on a single glance at a display of the ground motion at just one of the sensors. To make this decision the analyst compares the observed signal with a mental library of signals from previous events searching for a match. This procedure alone is perhaps not too difficult to automate, but the analyst's ability to compensate for changes in background noise levels, for possible glitches in the signal itself, etc. are much harder to emulate. In order to build automatic systems which can match the performance of human analysts, we must seek to capture as much of this complex pattern recognition process as possible in the detection system.

Automatic event detection systems

The current generation of automatic event detectors (e.g. GSETT3/IDC -- Le Bras et al., 1994a; USGS/Pasadena -- Doug Given, pers. comm., November 1995; USGS/Menlo Park, Johnson et al., 1995) have several common features. All of them process waveforms to produce triggers for which some associated measurements are made (e.g. signal-to-noise-ratio (SNR), period, amplitude). Typically these triggers are

declared by processing the data using an algorithm based on the ratio of a short term data average to a long term data average (STA/LTA) and comparing the output to a pre-defined station dependent threshold. Once the triggers are found, they are used to form or "build" events. This is the process generally known as association. The idea is simple enough -- to try to find physical locations in space and time which will produce predicted arrivals that match the observed triggers to some prescribed level of accuracy -- but depending on the number of triggers, their measured characteristics, and the characteristics of the network of stations used, this can be a complex process and errors can be made in association. Even if a perfect association algorithm could be developed, however, its product would only be as good as the input, i.e. as the quality of triggers. None of these systems has the ability to "look" at the full waveforms as the trained analyst does; in each case the waveforms have been reduced to a series of discrete arrivals. Thus, there may be an inherent upper limit on performance of trigger-based systems because of their fundamental design.

These systems suffer from other problems as well. The association codes often are exceedingly complex because they try to map the complexity of the patterns seen by the trained analyst (e.g. the relative amplitudes of given phases) into a series of simple logic statements (e.g. Le Bras et al., 1994b). This can make the codes difficult to tune and maintain by anyone not familiar with the network and the details of the code. Also, most of the current generation of detection systems use only a portion of the energy radiated by an event (typically the first arrivals), and consequently they can miss events which are readily apparent to analysts who can identify later phases. In some cases where the first arrivals are marginal these other phases are the key to event detection (e.g., Pomeroy et al., 1982).

Waveform correlation event detection

A new type of event detection system based on full waveform correlations for a grid representing potential event locations was presented by Shearer in 1994. Shearer used the system to search through 10 years of LP (long-period) IDA (International Deployment of Accelerometers) data to look for slow or silent events (Shearer, 1994). The detector is a grid-based automatic system which examines continuous data streams for entire (processed) waveforms that match some portion of a master set of waveforms. Because this system is based on actual waveforms rather than derived information, it has the potential to avoid the problems inherent in trigger-based systems. The system was created specifically to detect events which were missed by other systems because they lacked distinct first arriving P waves. Shearer's detector did find several of these types of anomalous events, but it also missed many large events that the conventional systems had built, so its potential as a system for routine monitoring use was not clear. In this paper we present preliminary results of our project to develop and test a CTBT-quality waveform correlation event detection system. Specifically we focus on the problems presented by using broadband data, by using a finer grid than that used by Shearer, and by seeking to increase considerably the bulletin quality (i.e. to decrease the number of missed events and false alarms).

Development and Testing of the WCEDS Prototype

Shearer (1994) presents the basic scheme in detail, but we will briefly summarize it here. The method is based on a discretized search through time of all points in a grid representing the complete set of possible event locations. A simplified diagram is shown in Figure 1. Events are identified by matching phase characteristics in a profile of processed waveforms (with distance ordering determined by calculated distance from the candidate grid point location) with a predicted set of processed waveforms (master image). The master image itself is discretized or “binned” in distance, so that there are actually a finite number of distances against which a given waveform can be correlated. When a candidate grid point is to be evaluated, each station’s waveform is correlated with the column of the master image with the nearest corresponding distance:

$$C_j = \sum_{i=1}^{N_t} M_{ji} D_{ji} \quad (\text{EQ 1})$$

where C_j is the correlation (dot product) for the j th station, M_{ji} is the column of the master image at the distance for which the correlation is being tried, D_{ji} is the data for the j th station, and N_t is the number of points being correlated. The overall output, O , of the detector is the sum of the station correlations:

$$O = \sum_{j=1}^{N_s} C_j \quad (\text{EQ 2})$$

where N_s is the number of stations in the network.

The algorithm is compact and efficient because the complex knowledge needed to “recognize” an event is contained in the master image, not in the algorithm itself. Further, a level of complexity in coding and inter-process communication is avoided by eliminating the need for a separate signal detector because this system works directly from the waveforms (though they are processed, as we will discuss below). Because the detector is grid-based, and hence the full range of possible solutions is always tested, the logic of finding events is simple: an event is declared to have occurred at any grid point where the total correlation is greater than some pre-defined threshold. The system maps extremely well to distributed processing because of the inherent redundancy in a grid-based system; the detection task could easily be divided among several processors by dividing the grid into sub-grids.

Waveform pre-processing

We follow Shearer’s decision to process the waveforms before using them in the detector. We refer to this as “pre-processing” because it occurs prior to execution of the detector. For a number of reasons Shearer used an STA/LTA algorithm to process

the data prior to correlation. First, the processing can enhance phases relative to the background, particularly the later-arriving phases, and this enhancement in turn

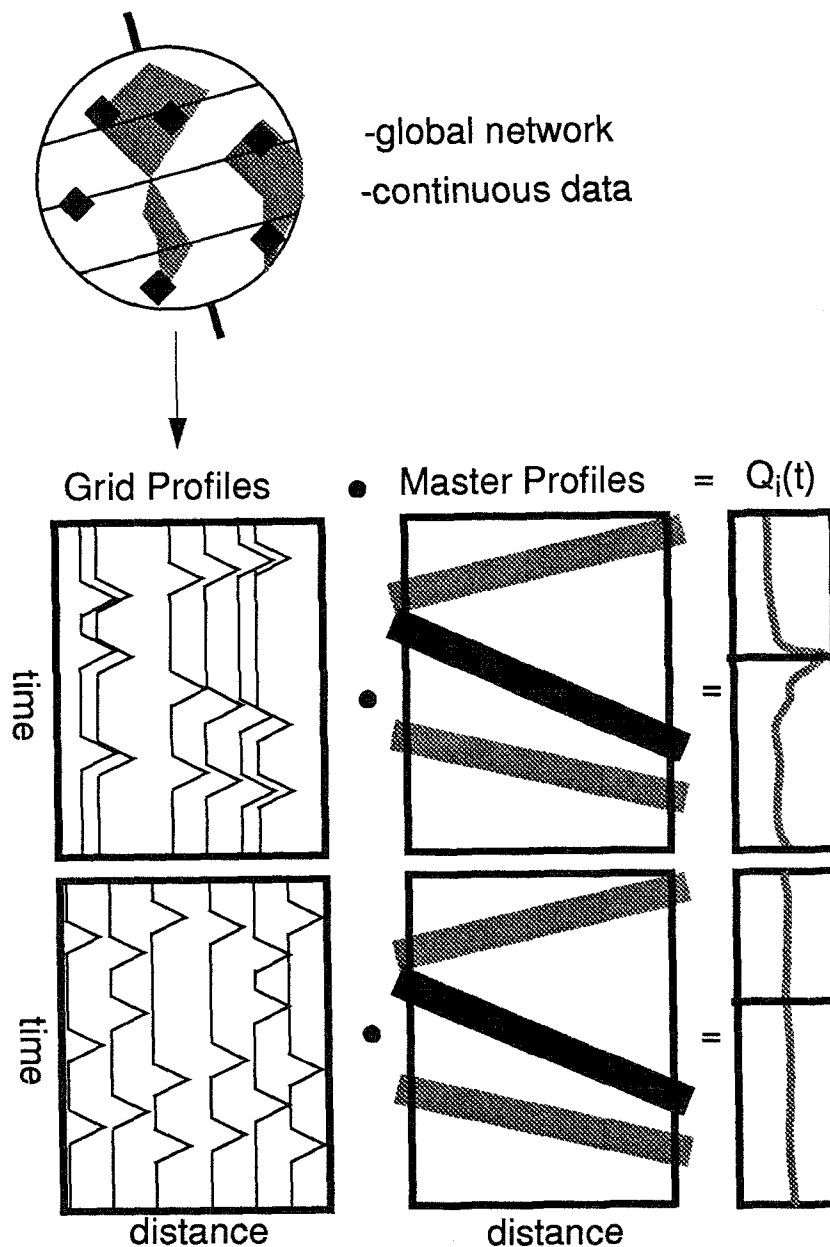


Figure 1. The basic waveform correlation event detection scheme

Continuous data from a global network is used to form profiles for each grid point which are checked against a master profile (image). A detection output time series, $Q_i(t)$, is generated for each grid point i . The variable shadings of the expected phases in the master profile indicate variable relative weighting.

should lead to lower detection thresholds. Second, the processing takes out the differences in amplitude due to instrument gain because the processed data is a ratio of data segments recorded by the same instrument. Without this normalization, the contribution of each station to the overall detection would be proportional to the gain of the instrument. Third, the processed data streams are smoothed; this makes it possible to use generalized correlation functions which significantly reduce the complexity (and number) of master images required (Figure 2).

There are many types of algorithms which could be used for pre-processing; for our system we favor methods which can balance optimal phase enhancement against minimal computation requirements (the latter factor is particularly important given the expected size of the CTBT network). For the prototype discussed in this paper a static STA/LTA was used, but given the importance of the pre-processing for detector performance, we have conducted an extensive review of various methods (see Withers et al., 1996). As a result of this analysis, we now believe that an adaptive STA/LTA algorithm (e.g. Tong, 1995) will yield superior results and we are currently working on modifying the system to work with data pre-processed in this manner. Note that while we have chosen to pre-process the waveforms, the code can operate on raw data. Creating a global waveform detection system based on raw data, however, is beyond the scope of our current efforts; it would require master images of a level of complexity that is inconsistent with currently available resources.

STA/LTA processing is computationally inexpensive and produces adequate phase enhancements, but we found that further processing is required for acceptable event detection. By definition, STA/LTA processing yields rectified (always positive) data streams and thus correlations with a rectified master image are also always positive. As a result, even random gaussian background noise will produce positive correlations and consequently the overall detection product for a hypothesized location will always be positive, making it difficult to determine if a small event has in fact occurred. Worse, if weights and/or number of phases in the master image are distance dependent and they generally will be, then the correlation product for processed noise will also be distance dependent. In fact, the output will be directly proportional to the area (integral) of the column of the master image against which the data is being correlated (refer to Equation 1). Due to this effect, we found cases where correlations for observable but weak signals would actually be less at some distances than for pure noise at distances with larger correlation areas. This would not be a problem if we used raw waveform data (demeaned and detrended) where the noise oscillates around a zero baseline so that correlations with a positive-only master image at any distance will be (on average) zero. Fortunately, we found that we could achieve noise cancellation with the processed data by removing the mean after the STA/LTA to once again yield a two-side (i.e. plus and minus) data stream. When this type of data was correlated with the master image, we found that noise has, on average, zero contribution, and we significantly improved the sensitivity of the system to small events.




Figure 2. Data pre-processing and waveform correlation

Correlating with raw waveforms implies a “wiggle-for-wiggle” match which will be station, phase, distance, and azimuth dependent. Correlating with a processed version of the waveform (squaring, fixed STA/LTA) is much simpler, and a generalized correlation function (in this case a boxcar) can be used.

Master images

As the above discussion should suggest, the master image is a critical part of the detection system. The master image contains the expected patterns of arrivals (and, potentially, other characteristics such as frequency content and directional information) in the data whose presence will indicate that an event has occurred. Referring to Equations 1 and 2, we can see that the output of the detector for a given grid location and time is proportional to the product of the data with the master image. Thus, the master image determines the relative weights of the data contributions. If it is

desirable to change the relative contribution of a particular phase/distance combination to event detection, this can be controlled through manipulation of the arrival amplitudes and arrival widths in the master image.

For Shearer's study, the master image was an empirical stack of waveforms from events with known locations, processed with the same STA/LTA algorithm as the data itself. This type of master image, which we shall hereafter refer to as an empirical master image (Figure 3), has two obvious strengths. First, it contains all of the phases which are actually present in the data and only those phases. Thus if, for example, a standard phase was known to be absent in a given area (e.g. Lg due to blockage) making an empirical master image for data from that area would yield a master image without the phase, as desired. Second, the relative weightings of the phases in an empirical master image are directly derived from the observed strengths of the data. Hence, if an empirical master image is used for correlation, weakly -observed phases will contribute very little to the overall detector output while well-observed phases will contribute a great deal.

These same features can lead to drawbacks, however. Beyond choosing the pre-processing algorithm and the network of stations used for data in the generation of an empirical master image, there is no convenient way to control the types of phases present in the image and their relative weightings. For example, if it were desirable to build an empirical master image-based detection system which would detect events based only on the P and S phases, generation of a suitable empirical master image would be difficult. Either a pre-processing scheme would have to be used which preferentially accentuated only those phases or a more general image could be generated and then carefully edited to eliminate all other phases. Further, even if the selection of phases could be controlled, the relative weightings could not be; they would be whatever was produced by the pre-processing and stacking. Using the STA/LTA scheme in its simplest form to generate a master image will lead to greater weights for closer phases, for phases with shorter path lengths, for phases from larger events, for phases observed at stations with lower noise levels, etc. Some of these effects may be desirable, but the overall effect is complex and difficult to control.

For these reasons we have chosen thus far to use master images based on travel time curves, as suggested by Shearer (1994). The process of generating a travel time-based master image is shown in Figure 4. One must choose the desired set of travel time curves, assign weights to the phases (possibly as a function of distance), and specify a spreading/shaping function. The strengths and weaknesses of travel time-based master images are nearly perfectly opposed to those of the empirical master images. With a travel time-based master image one has complete control of the choice of phases and their relative weightings, making it very easy to design custom master images such as for the P and S case discussed above. Also, variable weighting schemes based on phase type and distance can be easily implemented.




Figure 3. Empirical master images

This empirical master image was created from IRIS broadband vertical component data for events (1988-1994) with $M_w > 5.8$ and depth < 50 km. The data were processed with a static STA/LTA algorithm. See Shearer (1991) for details on forming empirical master images.




Figure 4. Travel time-based master images

(a) The master image is created by "smearing" a set of travel time curves. (b) The "smearing" is accomplished by convolution with a spreading function. Note the enhanced relative weighting of the second phase relative to the first and the third. Arbitrarily complex weighting schemes can be set up in this manner.

On the other hand, unlike an empirical master image it is necessary to completely specify which phases will be used and their relative weighting, and this requires additional information which may not be available. Without this information, one could specify phases which are not observed or whose observed timing is poorly matched by the travel time curves being used. Either of these effects would lead to degradation in the performance of the detector. Nevertheless, for the purposes of code development and testing we felt that the ability to control phase choice and weighting makes travel time-based master images a better choice and so far we have used them exclusively.

With either type of master image, the width of the phases used in the master image must be wide enough to: 1) match the width of the processed phases in the

data, 2) compensate for limited origin time sampling, 3) compensate for limited origin depth sampling, and 4) compensate for inaccuracies in the master image distance against which the waveform for a given station is correlated (due to both the discrete sampling provided by the grid and to the discretization interval of the master image). Proper widths for processed phases can be determined by observation of processed event data, and are in fact dependent on the choice of pre-processing algorithm (hence additional complexity is involved in using adaptive STA/LTA which can produce processed signals with very different widths). For static STA/LTA processing, signal pulse widths are fairly uniform and a single width can be used. Origin time and depth effects are also straightforward. To correct for origin time discretization, one simply widens all intervals by half of the time discretization at which the correlations are to be calculated. The depth discretization is a more complicated effect but can be approximated by a similar uniform widening dependent on the depth range to be spanned by each master image.

The distance discretization effects require a more complex phase and distance dependent adjustment of the correlation widths (Figure 5). For a uniformly spaced grid of spacing L the largest possible distance shift, Y , is controlled by the grid spacing (Figure 5b):

$$Y = L/(\sqrt{2}) \quad (\text{EQ 3})$$

On top of this one must apply the effect of the distance discretization in the master image; all grid point to station distances must be rounded to the nearest distance represented in the master image before correlation. To understand these effects, consider an example where a station is 72.6 degrees from an event and the grid spacing is 2 degrees ($Y \sim 1.4$). The possible range of distances to the nearest grid point are 71.2 degrees to 74.0 degrees. However, if the discretization interval of the master image is 1 degree, then the range of distances to the nearest grid point for purposes of correlation is 71 degrees to 74 degrees. Given a set of travel time curves, the grid spacing, and the master image distance discretization, one can generate the required correlation widths for stations at every distance. The process is illustrated in Figure 5c. Note that the width of the correlation interval at a given distance is a function of the slowness (slope) of the travel time curve at that distance: the greater the slowness of the phase, the wider the interval.

Event detection

In Shearer's original long-period study the operation of the system was essentially as diagrammed in Figure 1. For each origin time to be tested (2 minute increments), IDA data profiles (~ 20 stations) of a specified correlation length (3 hours) were created for each grid point (10 degree spacing for the Earth's surface only, leading to 416 points) and then correlated (see Equations 1 and 2) with an empirical master image to generate an output at each grid point. Over time, this led to a times series for each grid point which could be compared with an event detection threshold to determine whether an event occurred at a given grid/time point. Because of the





Figure 5. Master image correlation widths and distance discretization

(a) The actual correlation width must be wide enough to compensate for various sources of inaccuracy including the discretizations of the grid and master image. Note that the necessary discretization width is related to the slope (slowness) of the phase: the greater the slope of the phase, the wider the discretization width. (b) Assuming a grid spacing of L , the circles show the maximum distance (Y) which an event could fall from a given grid point without entering the domain of another grid point: $Y = L/(\sqrt{2})$. (c) Arrow #1 shows the possible shift in distance due to grid discretization. Arrow #2 shows the rounding of the grid shift to match the discretization of the master image (here 1 degree). The shaded bars show the width of the correlation function needed at each distance to compensate for the combined effects.

coarseness of the origin time increment, the small number of stations, and the small number of grid points, it was possible to process 11 years of continuous data on a desktop workstation in less than two weeks. The quality of the output was mediocre in terms of number of events detected (4061 total, 65% of cataloged events with $m_b > 5.5$), but as a high-quality event bulletin was not the goal of the study, this was not unexpected.

Applying the technique to data from a broadband global network with the intent of achieving sufficient sensitivity to compete with or exceed the performance of existing global detection systems (e.g. Le Bras et al., 1994a) implies developing a considerably more sophisticated version of the system. For development and testing we chose a data set from the Incorporated Research Institutions for Seismology (IRIS) broadband network, which consists of ~60 three-component stations which record data at 20 sps. We pre-processed our data at the original sample rate to preserve high-frequency signals, but then decimated to 1 sps before correlation to reduce the number of computations needed for the correlations. 1 sps is probably a conservative time spacing given that we have seen no pulse widths in the pre-processed data (using a static STA/LTA) that are less than 5 seconds. For our initial development and testing we used an approximately uniformly spaced 2 degree surface grid, leading to a total of about 9,000 points. For reference, a similar grid with 1 degree spacing would have about 41,000 points.

For a production-quality system, subsurface grid points must be monitored as well, which could lead to a very large total number of grid points. Fortunately, complete geographic coverage is not necessary at depth, because man-made events are confined to the very near-surface and deep- or intermediate-focus earthquakes are observed only in very limited, well-defined regions (e.g. subduction zones). If we assume that future sub-surface events will occur only in regions where past events have occurred, we can considerably reduce the number of additional points which must be monitored. For example, if we elect to use 5 sub-surface grids to span down to the deepest recorded seismicity (about 800 km) for a 1 degree grid spacing and use full sub-surface grids this would lead to a total of about 200,000 additional grid points. However, if we cross-reference these grid points against 10 years of PDE events (1983-1993), we find that only about 9,000 additional points are needed to monitor areas where seismic events are known to have occurred. Nonetheless, the extension to depth is non-trivial because each depth will require a separate master image and hence separate correlations.

If we confine our analysis to the surface and use a 2 degree grid and a 1 second time discretization and implement Shearer's system design, the number of computations needed is still enormous. For each candidate epicenter if we use a 1 hour correlation length, the number of multiplications and additions needed for the correlations (see Eq. 1) is about 1.94 billion (60 stations \times 3,600 time points \times 9,000 grid points). Processing 2 hours of continuous data, testing origin times every second was projected to take about 27 days on a single cpu 50 MHz Sparc20 (the run was never actually completed). While the run time could certainly be reduced by distrib-


uting the work load over several processors, a significant gain can also be made by re-thinking the algorithm.

If we note that the master image has a distance discretization interval associated with it (i.e. it does not have infinite spatial resolution) then it can be deduced that for a particular origin time the waveform from a given station can only have as many unique correlations with the master image as there are unique columns (distances) in the master image (this is only true if the master image is azimuthally invariant, however, as is the case here). For example, if the master image spans distances from 0 to 180 degrees with a distance interval of 1 degree, then for a given origin time the station can only have 180 possible correlations with the master image. To put it another way, there are only 180 unique distances that the station can be correlated with regardless of how many different grid points (i.e. candidate event locations) are monitored. Thus, even if we assume that every master image correlation distance is needed to monitor the grid points, the total number of additions and multiplications that are necessary to monitor a 2 degree grid for this master image is: 60 stations x 3,600 time points x 180 distance points = 38.9 million. For a 1 degree grid spacing we get the same number because the number of possible unique correlations does not depend on the number of grid points; it depends only on the distance discretization of the master image. Comparing these figures with those calculated above, it is clear that there is a tremendous amount of redundancy in the correlations (see Figure 6a), and the redundancy grows as the resolution of the grid increases (assuming that the discretization interval of the master image remains the same). If the system is redesigned to eliminate the redundancies in the correlations, the reduction in the number of multiplications and additions needed to calculate the correlations to monitor a 2 degree grid is a factor of ~50, and to monitor a 1 degree grid it is a factor of ~220. The difference between the number of multiplications and additions for the original method versus the new method is that the former multiplies by the number of grid points while the latter multiplies by the number of discrete distances in the master image. Thus the new method should become increasingly efficient as the number of grid points increases relative to the number of distance bins in the master image. If the number of grid points is smaller than the number of master image distances, then the original method will be more efficient. In the Shearer study the number of grid points was 416 while the number of distance bins was 360 so there would have been very little improvement with the new method.

A redesigned system taking advantage of the idea of computing all of the unique correlations upfront is shown in Figure 6b. In this scheme, the master image is loaded into one matrix (M), with each row defining a distinct distance, while the data from each of the stations is loaded into the columns of another matrix (D). A full matrix product is computed between the two which yields a correlation matrix (C) which contains all of the possible unique correlations for a particular origin time.

$$C_{ij} = \sum_{k=1}^{N_i} M_{ik} D_{kj} \quad (\text{EQ 4})$$

(a)



(b)

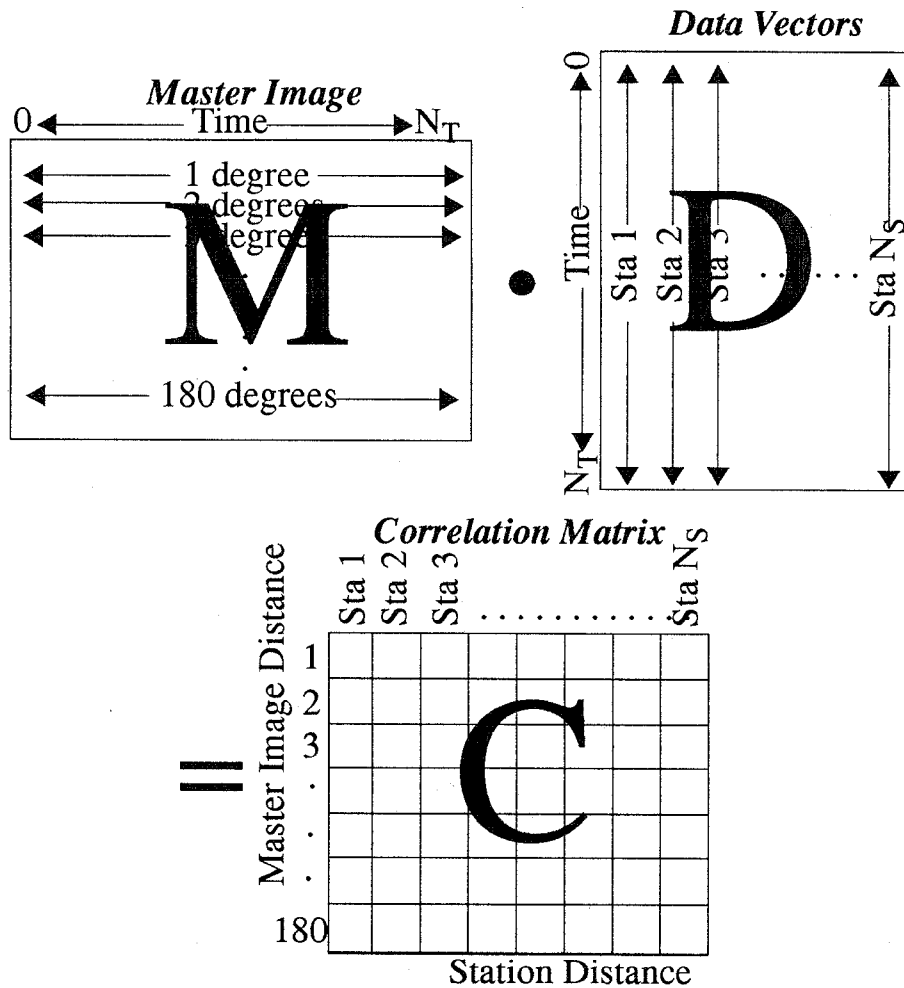


Figure 6. Correlation redundancy and a new detection algorithm

(a) The original code is inefficient for a large number of grid points because even though many grid points lie the same distance (r) from a given station and therefore will have the same dot product contribution for that station, new dot products are always calculated for each grid point-station pair. (b) The problem can be re-cast more efficiently as a full matrix multiplication of a data matrix (D) and a master image matrix (M).

The detector output at each grid point is determined by summing elements in C:

$$O = \sum_{j=1}^{N_s} C_{ij} \quad (\text{EQ 5})$$

This is essentially the same as Equation 2 except that row index (i) of C_{ij} is now a function of the grid point being checked. Calculating the detector output for each grid point implies following a particular summation path through the columns of C. A very efficient algorithm can be set up if, for all grid points, one pre-computes the distances to each station rounded to the nearest discrete distance in the master image and stores these before the detector is run. These distances then define the row in each column of the C matrix which are needed to calculate the detector output for a given grid point. The savings in computation using the new method can be tremendous for a large grid. The projected 27 day run using the old method discussed above was completed in just 7 hours using the new method on the same machine.

The correlation matrix

Once calculated, the C matrix contains the correlation information needed to monitor every grid point for a particular origin time (when the origin time advances, the D matrix changes and so C must be recomputed). The individual cells represent the correlation of an entire observed waveform with a waveform from the master image. All observed phases which are present in both the data and the master image will contribute to the value: a poor correlation could represent the lack of an observed phase in the data, the lack of a corresponding phase in the master image, or both. Each column of the C matrix represents the set of all of the correlations for a given station for the full range of distances in the master image while each row represents the set of all of the correlations of a given distance in the master image with all stations. The greater the value in a given cell the higher the correlation between the corresponding observed waveform and the given column of the master image and hence the greater the chance that an event occurred at the hypothesized distance from the particular station at the particular time. If an event did occur at the origin time for which C has been calculated then there will be high cell values at the correct distances for stations which recorded the event. If we re-arrange the columns of the C matrix (the station correlations) such that they are at the correct distances from the assumed epicenter and pad with zeros columns for distances where we have no data, then the correlations will align along the diagonal. In this ordering the horizontal axis can now be thought of as true distance while the vertical axis is correlation distance. The detector output for the assumed grid point is just the sum of the cells along the diagonal.

Figure 7a shows a C matrix calculated for the trial origin time 08:42:32 on October 2, 1993, sorted for the nearest grid point (a 2 degree grid spacing was used) to the epicenter of a m_b 6.2 event which occurred in the southern Xianjiang Province, China at 08:42:32.7. This is the largest event listed in the PDE (Preliminary

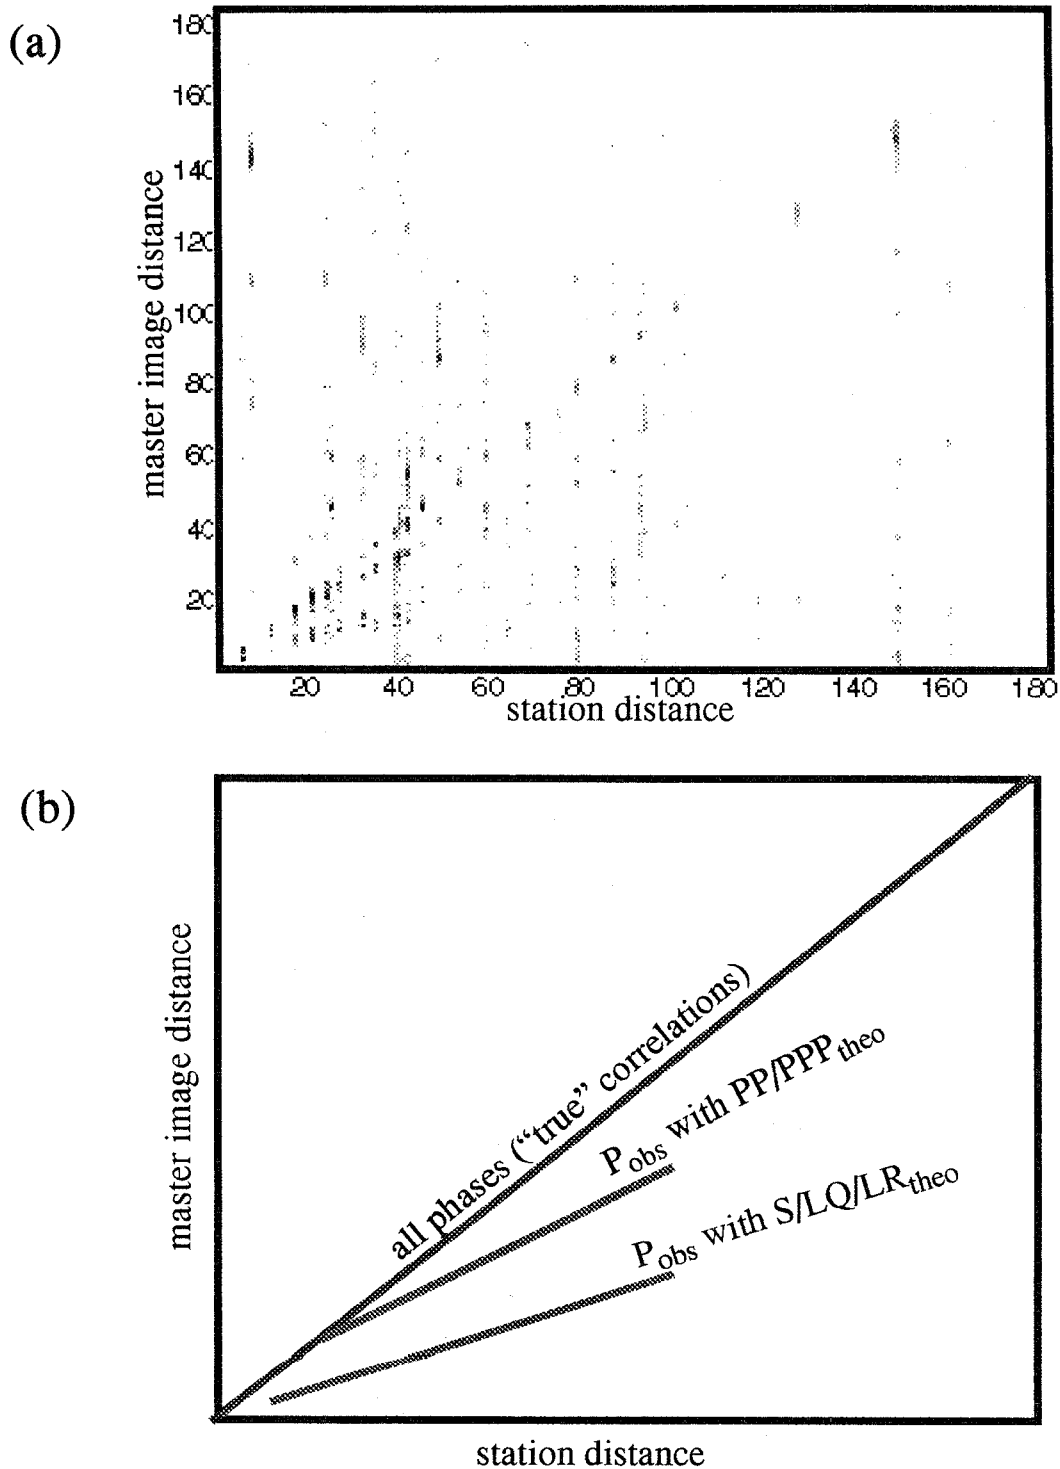


Figure 7. C matrix examples

(a) The sorted C matrix (extra zero columns were added to "diagonalize" the matrix) for a mb 6.2 event in southern China which occurred at 08:42:32.7 on October 2, 1993. (b) Our simplified interpretation of the C matrix: the three dominant features are the true correlation and two false correlations (sidelobes).

Determination of Epicenters bulletin, available from the National Earthquake Information Center) for several hours (see Table 1).

Table 1. PDE events for October 2, 1993 08:30:00 to 09:30:00

Origin Time	Latitude	Longitude	Mag.	Location
08 30 43.8	35.400 N	117.823 W	2.8	CENTRAL CALIFORNIA
08 30 53.0	60.406 N	4.951 E	1.8	SOUTHERN NORWAY
08 42 32.7	38.190 N	88.663 E	6.2	SOUTHERN XINJIANG, CHINA
09 00 13.6	53.049 N	158.611 E	5.0	E. COAST OF KAMCHATKA
09 20 12.2	38.206 N	89.284 E	4.9	SOUTHERN XINJIANG, CHINA

As mentioned above, our data comes from the three-component broadband IRIS network. For each station, the three channels were pre-processed separately (0.5-5.0 Hz 3 pole butterworth band pass filter, squaring to rectify and enhance coda phases, static STA/LTA with window lengths of 3 and 27 seconds) and then summed before correlation. The master image was derived from the IASPEI (1991) tables (0-3600 seconds, 1 second discretization; 1-180 degrees, 1 degree discretization; correlation function = 14 second wide sine function; phases = P, PP, PPP, S, SS, SSS, Lg, PKP). The predicted diagonal structure is intermittent (because many distances do not have corresponding stations) but clear, as would be expected for this event which was observable at nearly every station in the network.

Using the matrix formulation and a 1 second origin time discretization, we processed the interval from 08:30:00 to 09:30:00 and found that the maximum output of the detector occurred at the nearest grid point, within a second of the true origin time. The success of this test suggest that the detector can find large events with many correlations, but of course this is only a small part of the complete problem. The next step tested the ability of the system to build smaller events within the interval (see Table 1: all but the Norway event are visible in the IRIS data). To do this it was necessary to understand the problem of false correlations and to develop a methodology to deal with them.

False correlations

Returning to Figure 7a, we can see that in addition to the expected line of correlations along the diagonal, there are other distinct alignments with high correlation. The values along these lines are not as great as those along the diagonal, but are still much greater than the background. These other lines are the correlations of

observed phases with incorrect theoretical phases, i.e. the sidelobes of the correlations. In particular, we can pick out the correlation of Pobs with PP/PPPtheo (notice that they begin to split near 40 degrees where PP and PPP begin to separate), and Pobs with Sttheo (Figure 7b). We will refer to these correlations as "false" correlations to distinguish them from the "true" correlations which contribute to the detection of a real event. To understand how false correlations occur, consider a simple hypothetical system with three phases: A, B, and C. For a data set we will use a waveform for one station at which all three phases are observable and distinct. Let us first consider correlations for the correct origin time (Figure 8a). To create the column in the correlation matrix for this station, the observed waveform will be correlated with all of the distances represented in the master image. This will lead to a proper correlation of all three observed phases with the corresponding theoretical phases at D_3 , but also to two false correlations at D_1 and D_2 where observed phases align with theoretical phases which do not correspond. The true correlation is always as big or bigger than the others because it will have at least as many contributing phases and often more, but the false correlations may still be quite large and this can be a problem. For grid points at distances of D_1 or D_2 from this station, the false correlations

will contribute to the detector output. If enough false correlations of sufficiently high values lie at the appropriate distances for a given grid point, the summation may exceed the detection threshold and a false event could be declared.

Let us consider more carefully the effect of true and false correlations on the output of the detector. In the three phase system discussed above, the true correlation at D_3 will lead to high output on a ring of grid points at a distance of D_3 around the station (the output will be a ring because no azimuthal information is used); the grid point nearest to the true location should lie somewhere on the ring. Without additional stations or additional information, however it would be impossible to choose the correct point. The false correlations will also define high output rings of grid points, one with radius D_1 and the other with radius D_2 . As discussed above, the D_3 ring should always have values as great or greater than the others, but the grid points on the other rings will have detector outputs well above the background levels and may be greater than the event declaration threshold.

In addition to the false correlations at incorrect distances for the correct origin time, we will also have false correlations for erroneous origin times at the same distance (Figure 8b). Again, the largest correlation will occur at the correct origin time because only at this time will all phases contribute, but other correlations will occur at other origin times and these may also lead to falsely declared events. The actual case is a generalization of the two simplified ones shown; each waveform will falsely correlate for true and false distances over a range of origin times spanning from before till after the correct origin time. Thus when an event occurs, its correlation effects are smeared over both space and time, and as a result other events that occur within the range of the time smearing may be obscured.

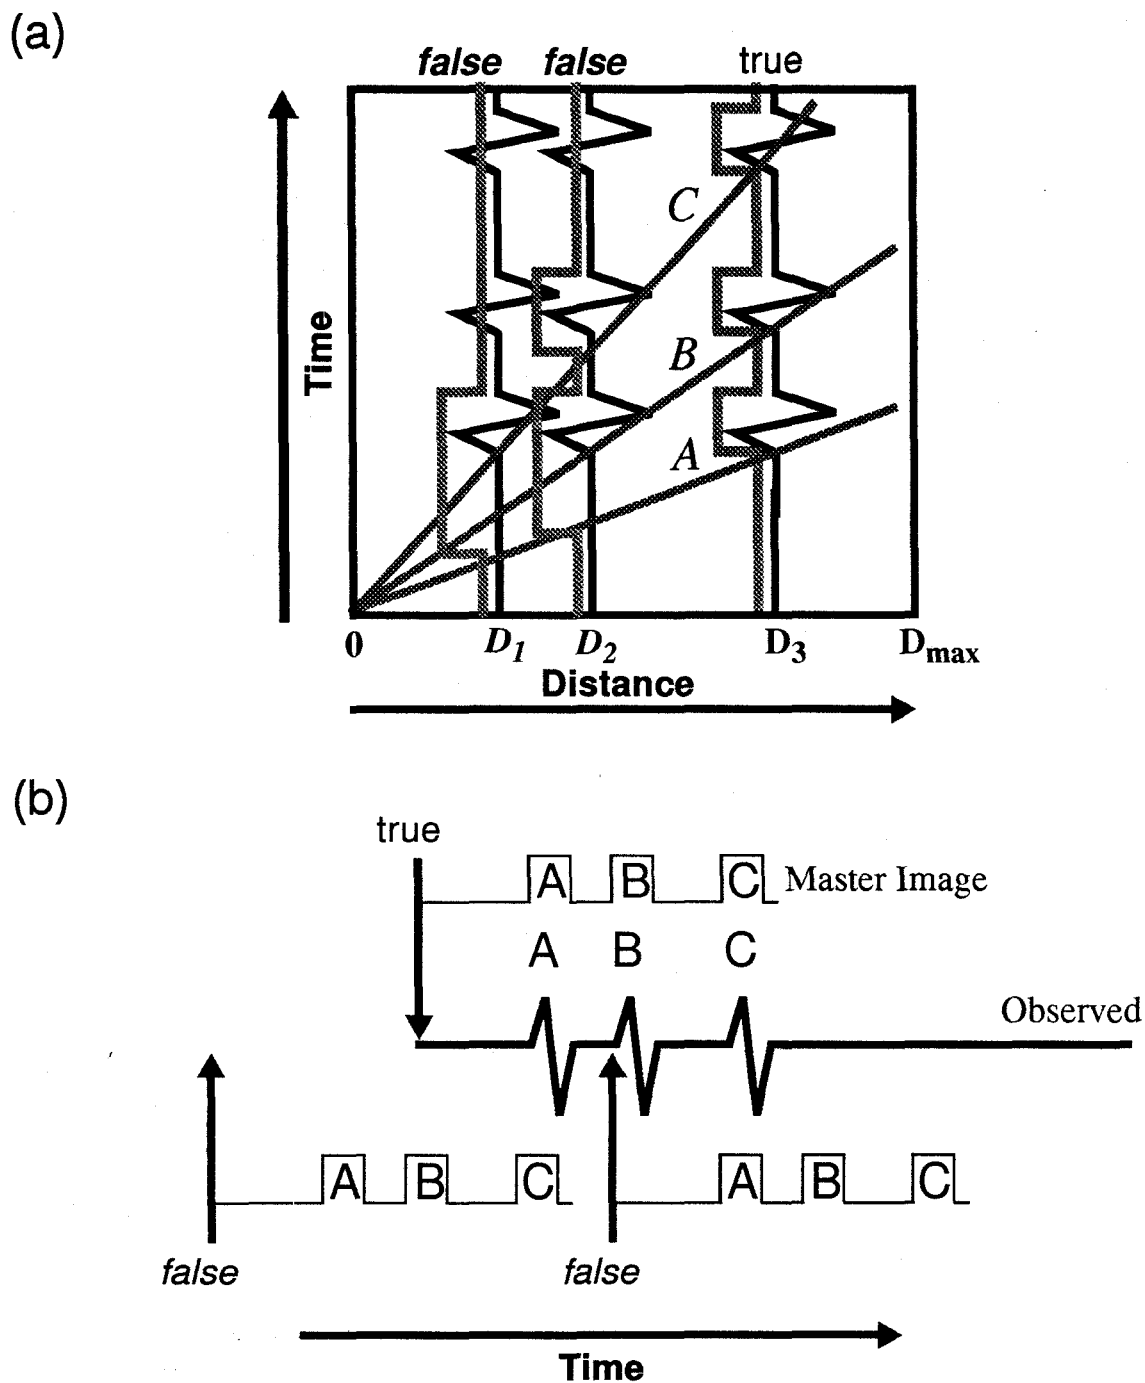


Figure 8. The generation of false correlations

A demonstration of how false correlations occur using a simple three phase system. (a) For the correct origin time the greatest correlation will occur at the correct distance, D_3 (A-A, B-B, C-C), but lesser correlations will occur at D_1 (A-B, B-C) and D_2 (A-C) as well. (b) At other origin times other false correlations will occur also.

Event stripping and event masking

A simplified hypothetical case of a time interval with three events having intermingled phases is shown in Figure 9a. Consider how our detection system would fare on this interval. The largest event (#2) should be easily found by searching for the maximum overall output within the interval. Building the other events, however, will be problematic. If, instead of searching for the overall maximum, we search for all time/grid point combinations which exceed a threshold, we might be able to find events #1 and #3, but we would also declare a huge number of false events in the process. To detect all three events without additional false events, the events must be processed sequentially, with all effects of the largest detected event removed from the data before searching for the next event. Figures 9b and 9c show how this works. As each event is detected, its effects (i.e. true and false correlations) are stripped away and the detector output is recalculated. At each stage in the process, an event is found by searching for the overall maximum. The process can be repeated until there is no output above the correlation threshold.

This process mimics the standard procedures of many human analysts and that it is present in some form in virtually every seismic detection system which must deal

with events whose phases can overlap in time. An experienced analyst who wants to try to find all of the events in a given time period in the most efficient manner will start with the largest event; this is true whether the analyst begins with raw data or with a set of events which have been built by some automatic system. The largest event is selected first because it is to be expected that this event will account for the most observable phases in the time period. It is essential to correctly associate all observable phases from this event because any remaining phases are candidates for building other events. Trigger-based automatic association systems prioritize in the same way. The association algorithm is handed a pool of triggers corresponding to a given time segment to build events from. The algorithm will try to build the largest event first, perhaps by looking for the triggers with the largest SNR. Once a reliable epicenter has been determined (generally using first arrivals) the system will thoroughly examine all of the remaining triggers looking for possible later phases before attempting to build a second event.

There are several possible ways to accomplish event stripping in our system. The most obvious would zero the waveforms where phases have been observed for a detected event. This implies removing data, which tends to make seismologists uneasy, but it is really no different than what goes on in trigger-based systems. Virtually no system will allow the phase observed at a given time to be associated with more than one event: a choice must be made (this is the topic of "conflict resolution", e.g. Beall et al., 1995). Once a trigger has been associated with a given event it is effectively removed from further consideration (though more complex systems do allow re-checking) which is equivalent to zeroing the waveform from which it came.




Figure 9. Event stripping

The top panel shows the detector output for a hypothetical sequence of three events. Simply identifying events as peaks and subpeaks would not work well: we could detect event #2, and possibly #1, but we would get many false events as well and event #3 probably would not be detected at all. If we use an iterative event stripping scheme, however (shown in the sequence of panels, from top to bottom), events #1 and #3 can be detected without any false event detections.


We have opted for another approach, however, because in our system the censored waveforms would have to be re-correlated with the master image and this is a computationally expensive process. In the two-hour processing example discussed previously, we found that calculation of C took more than half of the total computation time involved in running the system. For this reason, we chose to mask the correlations in the C matrix that are associated with a detected event and then recompute the output (O ; Equation 5) at the grid points. We wish to note however, that as we have continued to refine the code we have greatly increased the speed at which the C matrix can be recomputed (in fact, full recomputation is not even necessary) and recomputation of C may now be viable. For the purposes of this paper, however, we will use the masking technique described below exclusively.

Returning to Figure 7, we can see that what we must do is mask the lines of true and false correlations before looking for any other events. To do this we set up another matrix of the same dimensions as C and use it to keep track of which cells of C have been masked. We will call this masking matrix the X matrix (Figure 10a). Re-running the detector, i.e. summing through the C_j as defined in Equation 2, using the masked version of C should allow us to build a smaller event if it is present. As discussed previously, the correlations occur for a range of times from before until after the true origin time, so the X matrices must be generated for the entire range of times and correlation matrices must be re-calculated for this range (Figure 10b).

To calculate the X matrices when an event has been declared, it is necessary to know which stations contributed to the event detection and which phases contributed at those stations. It is the combination of these observed phases with the phases that are present in the master image which leads to the correlations. We could simply assume that all stations contributed all possible phases for every event, but this would be unwise because masking cells in the C matrix implies that certain station distance combinations will be excluded from input to the detector and any of these could potentially be critical to the detection of another event. Hence it is essential that we accurately determine exactly which stations and phases contributed. Determining the contributing stations is straightforward. Referring to Figure 7, we can see that one can define a contributing station threshold value and test the diagonal value of each station against this threshold. In our current system we use the same threshold for each station, but it may be prudent to implement a more complex scheme wherein the threshold could vary by station and perhaps also by distance from the detected event.

Determining which observed phases are present is only slightly more complicated. Because an event has been declared at a given grid point, we know the column of the master image with which each station was correlated and we can use this information to determine candidate phase contribution windows. These windows can be used to check the data for contributing phases. The process is shown in Figure 11.

(a)



(b)




Figure 10. Use of the X matrix

(a) Once an event has been detected all of the correlations (true and false) are loaded into the X matrix which is then used to mask the C matrix to check for other events.

(b) Because the correlations occur for a range of origin times, X matrices must be generated and applied for a range of origin times.

Each time an event is detected, its correlation information is added to the set of X matrices which already have the information for any previous events.

So far we have opted to use a simple average value of the data within the candidate phase window compared to a threshold value to decide if a phase has contributed, but more sophisticated schemes (e.g. evaluation of total energy) could be implemented. Using these methods, we can quickly determine both the contributing stations and time windows of contributing phases and this information can be compared

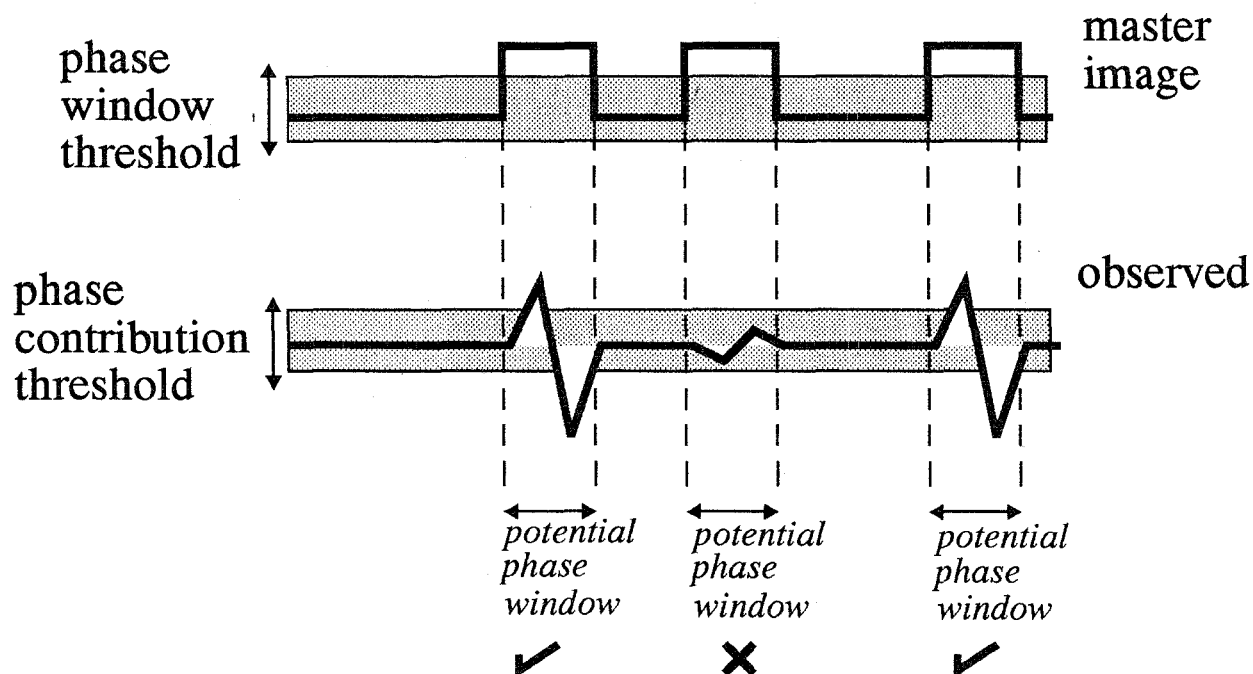


Figure 11. Determining the contributing phases

The master image can be used to define the windows in which to check for contributing phases. In this simple scheme, a phase is considered to have contributed if a phase contribution threshold is exceeded.

with the master image to set up the X matrices. For each contributing station and distance in the master image, we slide the time series of observed phase intervals past those of master image phases and note the origin times where the intervals overlap. The corresponding station-distance combinations in the X matrix for each origin time are filled in.

We used the same data set processed for Figure 7 to test the effectiveness of the X matrix. Figure 12 shows the detector output for a 2 degree surface grid on a global map before and after masking (the pre-processing and master image parameters were the same as used for the C matrix in Figure 7). Prior to masking, the maximum value occurs at the grid point nearest the epicenter in China. This maximum value occurs at the intersection of the true correlation rings of several stations, as expected. False correlation rings around several of the stations can also be seen, and it is apparent that some of these intersect as well and could easily lead to the detection of false events if we try to build other events without first removing the effects of the contributing phases of the detected event.

The masked output is shown in the lower panel. The absence of all correlation rings associated with the Chinese event (true and false) in the masked output map

demonstrates that the masking is working as intended. In this case there is no other event occurring at the same origin time so very little is left. If another event had occurred, however, its correlation structure would now be readily visible. A verification of this is provided by the presence in both maps of a correlation ring off the Pacific coast of South America due to a calibration pulse at the Chinese station LSA which correlated with PKP in the master image at a distance of 160 degrees. Because this observation had nothing to do with the event which was masked, its correlation structure remains after masking.

With the event masking procedure implemented, we have a system which should be able to process an interval of data and build events with intermingled phases. To test the augmented system we re-processed the October 2, 1993 hour interval using the same pre-processing, master image, and grid parameters. Events were built sequentially by seeking the grid/time point with the maximum detector output and then masking before repeating. The detector stopped when no grid/time point could be found which exceeded a minimum threshold. The results were very encouraging (Figure 13). Our system found a total of 5 events before falling below the chosen threshold, including all 4 of the observable events from the PDE. The order in which the events were found might appear random in that it does not follow the order of magnitudes, but it makes sense when one considers how much each contributes to the correlations in the C matrices (from which the detector output is produced). The 6.2 Chinese event had by far the most high correlations (it was observed by nearly every station in the network) and so was built first. The small California event was built next because of the large number of IRIS stations in California. There were 10 stations within 3 degrees of the epicenter and 8 of these had good observations. Hence despite its small magnitude, this event contributed significantly to the correlations in the C matrices and so was built second. The third and fourth events were both larger than the California event but had fewer close stations and hence fewer high correlations. The fifth event found was false, and an examination of it provides interesting information about how the system works. This event was created from un-masked surface waves of the 6.2 Chinese event. These phases could not have been masked because they were not included in the master image. Anticipating this problem, we had high-pass filtered the data to suppress LP signals, but apparently the filter selected was not completely effective. In principle, one could include surface waves in the master image and thereby mask for them if they are observed but in practice this is difficult because surface wave travel times are very path dependent and therefore the detector would have to use either an azimuthally dependent master image or a master image with correlation widths wide enough to span the range of possible arrival times. For our purposes the former option is presently unacceptable because it considerably increases the computational resources needed to run the detector while the latter must be rejected because the required masking would blind the detector to huge intervals of data where potentially interesting events might occur. Thus we prefer to try to eliminate surface waves by pre-processing, which may imply an occasional false event until better pre-processing can be implemented.

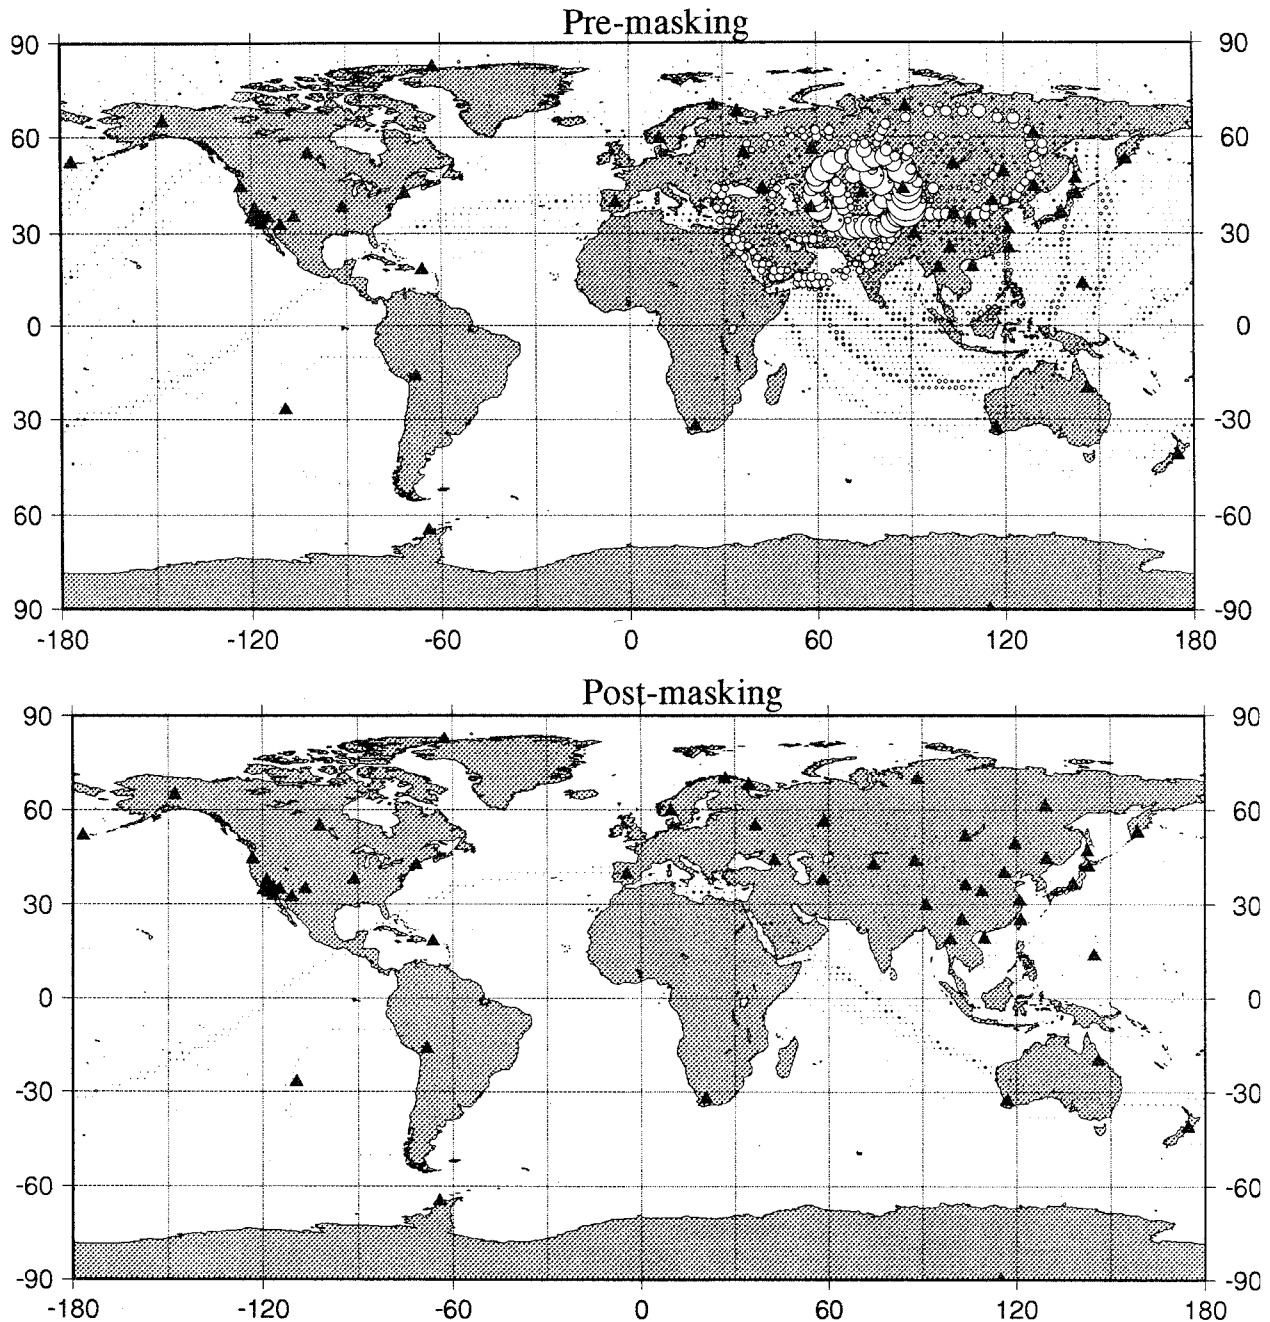
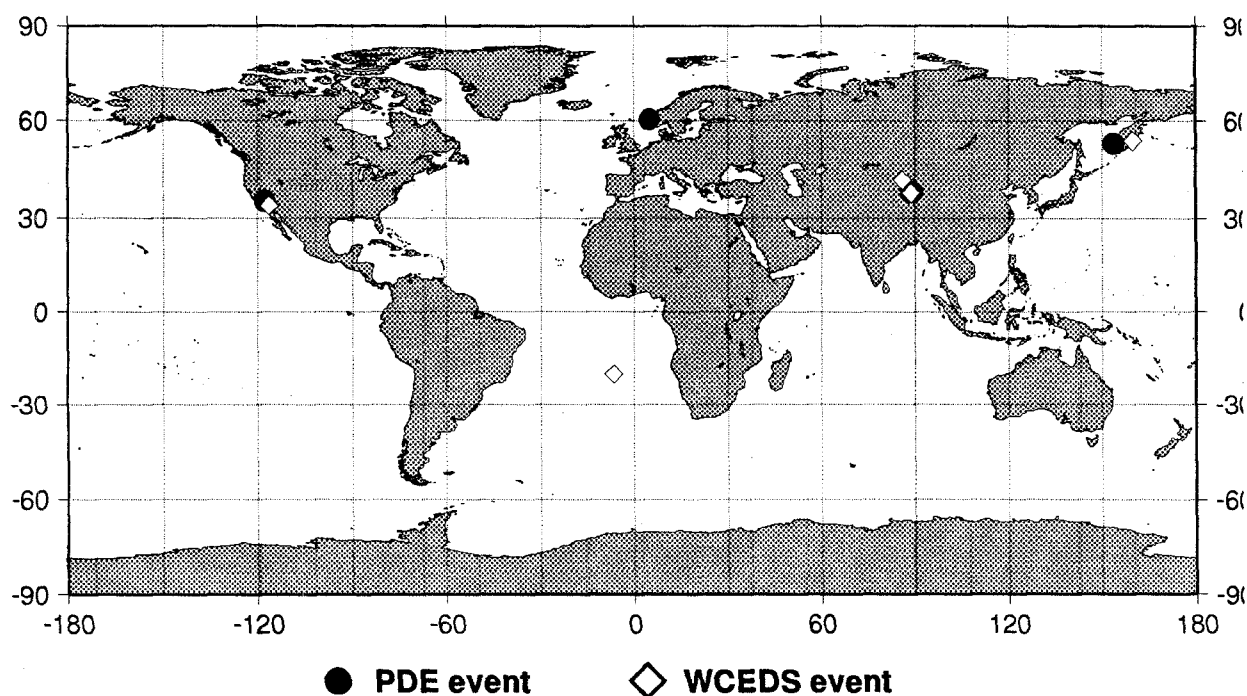


Figure 12. Pre-masking and post-masking detector output

The IRIS broadband stations are shown as dark triangles. The origin time for the calculation is 08:42:32 on 10/02/93. In each figure a circular symbol is plotted at each grid point with the radius of the symbol proportional to the detector output at that grid point. **(top panel)** Pre-masking output. The largest symbol occurs at the grid point nearest to the epicenter in China where the true correlation rings for several stations intersect. **(bottom panel)** Post-masking output. Nearly all of the correlation rings-- true and false -- are gone. The apparently stationless ring off the Pacific coast of South America is actually a 160 distance correlation ring from a calibration pulse at the Chinese station LSA.



PDE Bulletin (5 events)

(in order of occurrence)

- 2.8 California
- 1.8 Norway (not detectable)
- 6.2 China
- 5.0 Kamchatka
- 5.0 China

WCEDS Bulletin (5 events)

(in order of detection)

- 6.2 China
- 2.8 California
- 5.0 Kamchatka
- 5.0 China aftershock
- bogus - built from surface waves for China mainshock

Figure 13. WCEDS vs. PDE

The data come from the IRIS broadband network. The segment processed is 08:15 to 09:15 on 10/02/93.

Continuous execution

While the above results are encouraging, bulletin quality is not the only criterion which must be considered for a monitoring system: near real-time response is also essential. For this reason event detection systems used for continuous monitoring divide data into small segments to insure that events are detected as soon as possible: the minimum length of the segment which can be processed is constrained by the requirements of the detection system. In our case, to run the detector for even one potential origin time, we must have data from that origin time through the time span of the master image or all of the potentially available phases may not contribute to the correlation. Presumably systems that need very quick response times would have to use shorter master images.

Let us consider how the processing of segments could occur by considering an 8 hour interval of data with 4 events in it. This is shown in Figure 14. If we process the data as a single long interval (Figure 14a) then all of the events present will be built properly, but the time delay in declaring the events is too large, particularly for an event occurring early in the interval. An obvious solution is to divide the interval into smaller segments and process these as they become available; in this case we could choose four 2 hour segments (Figure 14b). The difficulty with such a system is caused by events which occur near the end of a segment, and this is a common occurrence for a global network when segment lengths are short. Such an event may have phases which correlate in more than one segment and consequently the event may not be built properly. In this case two of the real events have correlations in multiple segments and as a result we declare two extra false events. Unless there is some communication between the processing of the segments, it is not possible to recognize these as false events.

To deal with this problem it is necessary to use some sort of overlap or "look back". The process is illustrated in Figure 14c. The master image has a correlation time length associated with it (the number of rows in the master image defines the time span), so it is apparent that any events whose origin times fall more than that time length before the end of the current segment can be trusted. These events will have secondary peaks at other grid point/time points which will be examined as the detector runs through the current time segment, but this will not matter because the true grid point/time point will also be examined and preferred because it must have a greater value than any of the false correlations. Hence the true point will be found, and all of its correlations will be correctly masked so that the detector can look for smaller events. The difficulty comes when it is not possible to guarantee that if a better grid/time point exists (i.e. the true point if we have found a false event) the detector would be able to compare the two points and select the correct one. Any events in the segment following the trusted event segment (see Figure 14c) fall in this category. As soon as it is possible to correlate with any phase corresponding to an event whose true origin time lies outside of the segment currently being processed, events built must be labelled as suspect. Note however, that these events still must be built

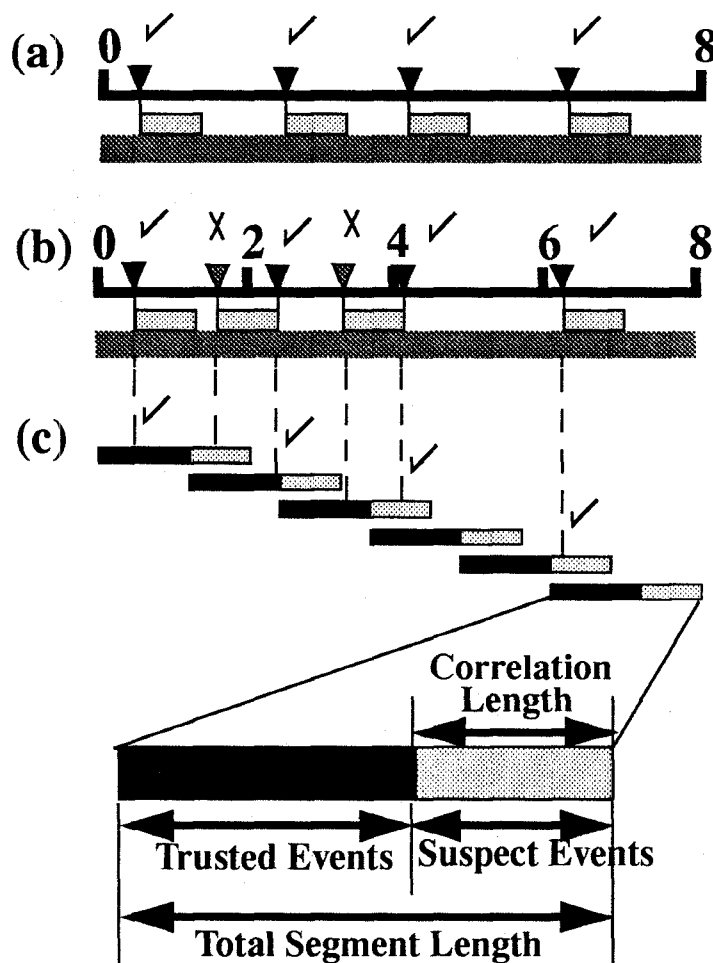


Figure 14. Continuous execution

Verification systems must be able to produce a bulletin in near-real time, which implies processing of consecutive segments of data. (a) Processing an 8 hour segment of data will produce the correct events (4 in this case), but requires an unacceptable delay in bulletin availability. (b) Breaking the 8 hours into two hour segments will meet the bulletin availability requirement but can lead to the creation of false events if the real events overlap the beginning and/or end of a segment. In this case, two false events are built. (c) The problem can be dealt with, however, by overlapping the 2 hour segments and introducing the concepts of trusted events and suspect events. The false events are built, but then discarded in favor of the correct events.

and masked because whether or not they are real, the correlations related to them may corrupt the association process in the current time period.

The key to continuous processing is to treat the trusted events and the suspect events differently: both the masking (i.e. the X matrices X_{trusted} and X_{suspect}) for each and the lists of events built must be kept separately. Except for this modification the processing is as discussed above: events are built and stripped away until no grid/time point in the segment exceeds the detection threshold. When the detector is ready to move on, the next segment to be processed must overlap the previous segment by the correlation time length of the master image, i.e. by the suspect event interval. As we begin processing, we start with whatever information is already in X_{trusted} , but we flush out X_{suspect} and the list of suspect events. Now we begin processing as above, except that what was the suspect interval in the previous segment is now part of what will be the trusted interval in the current segment. Thus, any of the flushed events which were legitimate will be rebuilt while false events will be rejected in favor of the true events. In this manner the process continues one segment at a time, perhaps building false events but correcting them before they are added to the final bulletin. We are currently in the process of implementing this scheme so that we can run our detector continuously.

Discussion

We have outlined an event detection system based on waveform correlation, presented examples of output, and shown how the system could be modified to operate in near real time. However, in order to meet the high performance requirements of some potential users (e.g. AFTAC, USGS) both the quality of the event bulletin and the speed with which it can be produced must be improved.

Some measure of speed improvement can be gained by using more powerful hardware and/or resorting to multiple processors if the basic design is amenable; ours is. The code itself is ansi C which could be run on a variety of machines with minor modifications, and we are in the process of rewriting the code in C++. With the exception of the event masking, nearly all of the processing that is done within the code could be distributed to an arbitrarily fine level. Pre-processing (filtering, STA/LTA, etc.) could be divided among workstations or processors within a single machine by station, channel, or time. The addition of new stations or changes to more computationally demanding pre-processing could be handled by using additional processors without additional delays in producing a bulletin. Similarly the computation of the C matrices and the calculation of the detector output at each grid point could be distributed in any number of ways as necessary. Thus, for organizations with large hardware budgets, almost regardless of the amount of data used it should be possible to configure a system which will be able to operate in real time.

Even if such resources are available, however, we feel that it would be prudent to pursue making the code more efficient and there are numerous opportunities for

this. Consider the way in which the detector finds an event in a time segment. In the current system, the output for each grid/time point is calculated by summing rows in the C matrix. The detector keeps track of the maximum value, and when all points have been evaluated, if the maximum exceeds the detection threshold an event is declared at the corresponding point, masking is set up, and the process is repeated. In each iteration we search the entire time period. If the segment has a high level of seismicity throughout, this method may be as good as any, but if the segment is characterized by isolated events, as seismic data often are, a large proportion of the C matrices will correspond to origin times when no events occur and a considerable amount of time is wasted in checking them. Significant efficiency improvements can be made if it is possible to determine whether the C matrix for a given origin time is likely to contain a hypocenter.

There are many possible ways to evaluate whether a C matrix is likely to have event information, but one of the simplest is to find the maximum correlation in each column (i.e. for each station) and sum these. We will call this value the maximum C sum. The maximum C sum represents the maximum possible output which the detector could have for the corresponding origin time, though in most cases it will not correspond to a real grid point (the grid points represent only a small subset of the total possible summation paths through the C matrix). This value must always be greater than or equal to actual detector output for the grid point nearest to a given event, thus it provides a convenient index to the presence of an event and it can be used to make our code more efficient. One very simple use is to compare the maximum C sum to the event detection threshold before computing any of the grid point summations: if the maximum C sum does not exceed the threshold then neither can any of the grid point summations and this origin time can be skipped. With this simple check we can avoid a number of summations equal to the number of grid points for every C matrix which cannot yield an event, and this can be a considerable savings for a large grid. The actual improvement will depend on the amount of seismicity and the threshold used, but in our 1 hour test case we were able to reduce the run time by about 1/3.

A more sophisticated and potentially much more powerful use of the maximum C sum is to identify the time sub-segment within which the event which contributes most to the C matrices is likely to lie and process only that sub-segment rather than the full segment. In this manner, the detector can quickly jump to a small time interval of interest potentially saving a huge amount of time. For high levels of seismicity this may be a critical feature of the system. In the current version, each time an event is found and the X matrices are loaded for masking, the code must go back to the first C matrix in the time segment and recalculate all of the detector outputs until the last C matrix is processed. Even with the max C sum check indicated above, this is a tremendously wasteful process. We have only recently begun to test the idea of using max C sum values to find the next event to detect and mask, but the results show promise: the maximum C sum values do seem to track events.

A considerable increase in the quality of the output (i.e. decreases in the number of false events built and real events missed) can be gained by building better phase recognition and noise rejection into the master images. It is the function of the master images to encapsulate the pattern recognition capability, i.e. to test the data for compatibility with an event hypothesis. In this paper the pattern recognition is fairly crude: the master images used require only that the processed data stream have amplitude within the specified time windows for the selected phases. While time and amplitude are important parameters to test for event existence, there are others which can also be used (e.g. azimuth, slowness, frequency content). In our future efforts we plan to add some of these features to the master images, increasing the level of sophistication of the pattern recognition in our system and thereby producing a better bulletin.

Another obvious means to improve the quality of the output is to use a better network of data. We chose to use the IRIS network for our preliminary testing because it is a global network, the data was relatively easy to obtain, and the pre-processing involved in using three component data is fairly simple. In the future we would like to test the system on the primary network used for GSETT3, which consists almost entirely of arrays and consequently should have much better sensitivity. To operate on the GSETT3 data we must develop pre-processing schemes which can take advantage of the signal enhancement properties of arrays without seriously compromising the performance of the system. We have just begun to work on this problem. Processing the GSETT3 data also gives us the opportunity to compare our results with those produced by the current event detection systems.

Conclusions

We have developed a working prototype of a waveform correlation event detection system which shows promise for treaty monitoring applications. The grid-based system is derived from an earlier long-period system developed by Shearer (1994) but has been greatly modified to suit the requirements of the monitoring environment. A tremendous increase in efficiency for large grids (greater than two orders of magnitude for a 1 degree grid spacing) has been realized by pre-computing all of the possible correlations for a given origin time and storing these in a correlation matrix (C) which can be used as a look-up table for computing the detector output. Using this system, events are found by searching each grid point/time in a given interval for the maximum detector output. Multiple events with intermingled phases can be detected by an iterative event masking process wherein the correlations (true and false) of each detected event are masked in C after the event is detected and the max grid point/time search is repeated using the masked version of C. This process is repeated until a cutoff threshold is reached.

The system has thus far only been tested on a small amount of data, but the results are encouraging. Running on an hour of IRIS broadband data for which 4 of the 5 events listed in the PDE were observable by the IRIS network, our system built

all 4 events. To test the system on larger data intervals we are in the process of implementing a scheme which will allow continuous operation by processing the data in overlapping intervals. With this scheme in place and another planned to allow us to effectively process array data, we will be able to process GSETT3 data and to directly compare our performance and bulletin quality to those from one of the most sophisticated of the current generation of automatic seismic event detection systems.

References

- Beall, G., R. Le Bras, W. Nagy, T. Sereno and H. Swanger (1995). Global association system phase 2: conflict resolution, *SAIC document #SAIC-95/1029*, 46 pp.
- Johnson, C. E., A. Bittenbinder, B. Bogaert, L. Dietz, and W. Kohler (1995). Earthworm: a flexible approach to seismic network processing, *IRIS Newsletter*, 14, 1-4.
- Le Bras, R., H. Swanger, T. Sereno, G. Beall, R. Jenkins, W. Nagy, and A. Henson (1994a). Global association final report, *SAIC document #SAIC-94/1155*, 28 pp.
- Le Bras, R., H. Swanger, T. Sereno, G. Beall, R. Jenkins, and W. Nagy (1994b). Global association design document and user's manual, *SAIC document #SAIC-94/1142*, 67 pp.
- Pomeroy, P. W., W. J. Best, and T. McEvilly (1982). Test ban treaty verification with regional data -- a review, *Bull. Seism. Soc. Am.*, 72, S89-S129.
- Shearer, P. M. (1994). Global seismic event detection using a matched filter on long-period seismograms, *J. Geophys. Res.*, 99, 13713-13725.
- Shearer, P. M. (1991). Imaging global body wave phases by stacking long-period seismograms, *J. Geophys. Res.*, 96, 20353-20364.
- Tong, C. (1995). Characterization of seismic phases -- an automatic analyzer for seismograms, *Geophys. J. Int.*, 123, 937-947.
- Withers, M., R. C. Aster, and C. J. Young (1996). Comparison of selected trigger algorithms for automated global seismic phase and event detection, *to be submitted to Bull. Seismol. Soc. Am.*

Distribution

- | | | | |
|---|---|---|--|
| 1 | Dr. Leslie Casey
U.S. Department of Energy
Office of Research and Development
Forrestal Bldg. NN-20
1000 Independence Ave. SW
Washington D. C. 20585 | 1 | Dr. Steve Bratt
Advanced Research Office
1400 Wilson Blvd.
Arlington, VA 22209 |
| 1 | Dr. David Russell (TTR)
Airforce Technical Applications
Center
Patrick AFB, FL 32925 | 1 | Dr. R. Blandford
Center for Monitoring Research
Suite 1450
1300 North 17th Street
Arlington, VA 22209 |
| 1 | Dr. Rick Schult (TTR)
Airforce Technical Applications
Center
Patrick AFB, FL 32925 | 1 | Jerry Carter
Center for Monitoring Research
Suite 1450
1300 North 17th Street
Arlington, VA 22209 |
| 1 | Dr. George Randall
MS-C335
Los Alamos National Laboratory
P.O. Box 1663
Los Alamos, NM 87545 | 1 | Tim Ahern
IRIS/DMC
1408 NE 45th St.
Seattle, WA 98105 |
| 1 | Dr. Dave Harris
MS L-205
Lawrence Livermore National Laboratory
P.O. Box 808
Livermore, CA 94550 | 1 | Bob Woodward
USGS-Albuquerque Seismological
Laboratory
Kirtland AFB, NM 87115 |
| 1 | Dr. Stan Ruppert
MS L-205
Lawrence Livermore National Laboratory
P.O. Box 808
Livermore, CA 94550 | 1 | Bob Hutt
USGS-Albuquerque Seismological
Laboratory
Kirtland AFB, NM 87115 |
| 1 | Dr. Dan Hagedorn
MS-KS-12
Pacific Northwest National Laboratory
P.O. Box 999
Richland, WA 99352 | 1 | Richard Aster
Department of Earth Sciences
New Mexico Institute of Mining
and Technology
Socorro, NM 87801 |
| | | 1 | Mitch Withers
Department of Earth Sciences
New Mexico Institute of Mining
and Technology
Socorro, NM 87801 |

1	MS0979	D. R. Breding, 5704
1	MS0979	D. B. Shuster, 5704
1	MS0979	L. S. Walker, 5704
1	MS0655	D. B. Carr, 5736
1	MS0655	E. P. Chael, 5736
1	MS0655	J. P. Claassen, 5736
1	MS0655	J. M. Harris, 5736
20	MS0750	C. J. Young, 6116
1	MS0750	M. C. Walck, 6116
1	MS0750	G. J. Elbring, 6116
1	MS0977	J. I. Beiriger, 9432
1	MS0977	R. G. Keyser, 9432
1	MS0977	B. N. Malm, 9432
1	MS0977	S. G. Moore, 9432
1	MS0977	J. R. Trujillo, 9432
1	MS9018	Central Tech Files, 8523-2
5	MS0899	Technical Library, 4414
2	MS0619	Review & Approval Desk, 12630 For DOE/OSTI